

Extended Dynamic Generalized Linear Models: the two-parameter exponential family

MARIANA A. O. SOUZA *

Universidade Federal Fluminense

HELIO S. MIGON †

Universidade Federal Rio de Janeiro

August 25, 2015

Abstract

We develop a Bayesian framework for estimation and prediction of dynamic models for observations from the two-parameter exponential family. Different link functions are introduced to model both the mean and the precision in the exponential family allowing the introduction of covariates and time series components. We explore conjugacy and analytical approximations under the class of partial specified models to keep the computation fast. The algorithm of West et al. (1985) is extended to cope with the two-parameter exponential family models. The methodological novelties are illustrated with two applications to real data. The first, considers unemployment rates in Brazil and the second some macroeconomic variables for the United Kingdom.

1 Introduction

Generalized linear models (GLMs) have become a standard class of models in the data analyst's toolbox. Proposed by Nelder and Wedderburn (1972), GLMs are widely used in many areas of knowledge. They allow modelling data of many different natures via the probabilistic description as an element of the exponential family and relating the response mean and the linear predictor in a non-linear form. The

**Address for correspondence:* Mariana Albi O. Souza, Departamento de Estatística, Instituto de Matemática e Estatística, Universidade Federal Fluminense, Rua Mario Santos Braga s/n, 7o. andar, Centro, Niterói, RJ, Brazil. CEP 24020-140. *E-mail:* mariana@im.uff.br.

†*E-mail:* migon@im.ufrj.br.

GLM class is a useful alternative for data analysis since it accommodates skewness and heteroskedasticity, besides allowing analysis using the data in their original scale. The evolution of these models as well as details regarding inference, fitting, model checking, etc., is documented in the seminal book of McCullagh and Nelder (1989) and many others in the recent literature.

The main criticism of the use of the one-parameter exponential family in certain applications is that samples are often found to be too heterogeneous to be explained by a one-parameter family of models in the sense that the implicit mean-variance relationship in such a family is not supported by the data. To overcome this limitation Gelfand and Dalal (1990) and Dey et al. (1997) introduced the class of two-parameter exponential family, which includes the ones presented by Efron (1986) and Lindsay (1986) as special cases. They argue that the introduction of a second parameter allows taking into account the over-dispersion usually present in the data, an issue that has been recognized by data analysts for many years.

During the 1990s, special attention was devoted to modelling the mean and the variance simultaneously. Taguchi type methods led to some efforts to jointly model the mean and the dispersion from designed experiments, avoiding the data transformation that is usually necessary to satisfy the assumptions of traditional linear models Nelder and Lee (2001). The process of quality improvement aims to minimize the product variation caused by different types of noise. Quality improvement must be implemented in the design stage via experiments to assess the sensitivity of different control factors that affect the variability and mean of the process. Nelder and Lee (2001) discussed how the main ideas of a GLM can be extended to analyse Taguchi's experiments. From a static point of view, the Bayesian inference for this class of models is fully discussed in the papers previously cited, while some alternative aspects of MCMC are discussed in Cepeda and Gamerman (2005) and Cepeda et al. (2011).

Our aim in this article is to extend the class of models introduced by Gelfand and Dalal (1990) and Dey et al. (1997) to deal with time series data and to propose a fast algorithm for estimation and prediction of this class of models. To reach this objective we propose an algorithm based on analytical approximations, for example, based on Laplace approximations. This way we are extending the conjugate updating method proposed in West et al. (1985).

The remainder of the manuscript is organized as follows. Section 2 introduces the class of models we are focused on. In Section 3 the conjugate updating of West et al. (1985) is extended to the two-parameter exponential family. Section 4 illustrates the proposed method with two case studies: the first one models unemployment rates in Brazil and the second one models some data on the UK economy as beta distributed data. Section 5 concludes with a discussion and possible future research directions.

2 Extended Dynamic Generalized Linear Models

In this section we introduced the class of extended dynamic generalized linear models (EDGLM). First we briefly revise the two-parameter exponential family and the dynamic generalized linear models, mainly aiming to fix the notation to be used in this paper. A special parametrization of the two-parameter exponential family is presented in this section. It is very useful to deal with data analysis when heterogeneity in the sample is greater than that explained by the variance function in the one-parameter exponential family. The distributions in this family are often used in many applications in the current literature, not only to deal with the topic of extra variability.

The two-parameter exponential family has the form

$$p(y|\theta, \phi) = a(y) \exp \{ \phi [\theta d_1(y) + d_2(y)] - \rho(\theta, \phi) \}, y \in \Upsilon \subset \mathbb{R} \quad (1)$$

where $a(\cdot)$ is a non-negative function, $d_1(\cdot)$ and $d_2(\cdot)$ are known real functions, $(\theta, \phi) \in \Theta \times \Phi \subseteq \mathbb{R} \times \mathbb{R}^+$ and $\exp\{-\rho(\theta, \phi)\} = \int a(y) \exp \{ \phi [\theta d_1(y) + d_2(y)] \} dy < \infty$. This is a suitable reparameterization of the general two-parameter exponential family as defined in Bernardo and Smith (1994).

This class includes many continuous distributions, such as the normal with unknown mean and variance, the inverse Gaussian and the beta distributions, parameterized by its mean and precision factor. The expression for the variances, as we will see in section 3.3, make clear the relevance of the precision parameter, ϕ , to control the model variance. Large values of ϕ corresponds to more precise data or equivalently with smaller variance. Some discrete distributions are also included in this class, such as the binomial (with the sample size known) and Poisson distributions, taking the scale parameter as fixed and equal to one.

Among other interesting features of this class of distributions, we stress the existence of a joint prior distribution for the parameters (θ, ϕ) in the form $p(\theta, \phi|\boldsymbol{\tau}) = \kappa(\boldsymbol{\tau}) \exp \{ \phi [\theta \tau_1 + \tau_2] - \tau_0 \rho(\theta, \phi) \}$, where $\boldsymbol{\tau} = (\tau_0, \tau_1, \tau_2)'$ and $\kappa(\boldsymbol{\tau})^{-1} = \int \int \exp \{ \phi [\theta \tau_1 + \tau_2] - \tau_0 \rho(\theta, \phi) \} d\theta d\phi$. Let $\boldsymbol{\psi} = (\theta, \phi) \in \Psi = \Theta \times \Phi$, to make the notation easier. Its prior mode and observed curvature matrix can be straightforwardly obtained differentiating the expression above with respect to the parameters vector $\boldsymbol{\psi}$. More specifically, the mode and curvature matrix satisfy the equations

$$\tilde{\boldsymbol{\psi}} = \arg \max_{\boldsymbol{\psi}} \frac{\partial}{\partial \boldsymbol{\psi}} \log(p(\boldsymbol{\psi}|\boldsymbol{\tau})) \quad \text{and} \quad J(\boldsymbol{\psi}) = - \frac{\partial^2}{\partial \boldsymbol{\psi}' \partial \boldsymbol{\psi}} \log(p(\boldsymbol{\psi}|\boldsymbol{\tau})).$$

Then it follows, after some algebra, that

$$\begin{pmatrix} \phi \tau_1 - \tau_0 \frac{\partial}{\partial \theta} \rho(\boldsymbol{\psi}) \\ \theta \tau_1 + \tau_2 - \tau_0 \frac{\partial}{\partial \phi} \rho(\boldsymbol{\psi}) \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \quad \text{and} \quad J(\boldsymbol{\psi}) = \begin{bmatrix} -\tau_0 \frac{\partial^2}{\partial \theta^2} \rho(\boldsymbol{\psi}) & \tau_1 - \tau_0 \frac{\partial^2}{\partial \theta \partial \phi} \rho(\boldsymbol{\psi}) \\ \tau_1 - \tau_0 \frac{\partial^2}{\partial \theta \partial \phi} \rho(\boldsymbol{\psi}) & -\tau_0 \frac{\partial^2}{\partial \phi^2} \rho(\boldsymbol{\psi}) \end{bmatrix}.$$

The predictive distribution is also defined in closed form, as

$$p(y|\boldsymbol{\tau}) = a(y) \frac{\kappa(\boldsymbol{\tau})}{\kappa(\boldsymbol{\tau}^*)}, y \in \Upsilon, \quad \text{where} \quad \boldsymbol{\tau}^* = (\tau_0 + 1, \tau_1 + d_1(y), \tau_2 + d_2(y))'. \quad (2)$$

Now that the basic notation is clearly stated, we can progress to the dynamic version of the extended generalized linear model. Let y_1, \dots, y_T be conditionally independent observations from the two-parameter exponential family and for each $t \in \{1, \dots, T\}$, denote $E[y_t|\boldsymbol{\psi}_t] = \mu_t$. Let us suppose that both the mean μ_t and the precision ϕ_t can be described by explanatory variables through possibly different non-linear link functions, denoted by g_1 and g_2 .

Therefore, given the prior moments of the latent states $\boldsymbol{\beta}_t$, the class of models to be considered in this paper is described by three components. The first is a conditional conjugate model describing observations in the two-parameter exponential family with its prior distribution:

$$y_t|\boldsymbol{\psi}_t \sim Ef(y_t|\boldsymbol{\psi}_t) \quad \text{and} \quad \boldsymbol{\psi}_t|D_{t-1} \sim CEf(\boldsymbol{\tau}_t), \quad \forall t = 1, \dots, T, \quad (3)$$

where $Ef(y_t|\boldsymbol{\psi}_t)$ denotes a distribution in the two-parameter exponential family (1), $CEf(\boldsymbol{\tau}_t)$ represents its conjugate prior distribution and D_{t-1} denotes all the information available up to time $t - 1$.

A general link function is introduced to relate the linear predictors with the mean and precision of the observational distribution evaluated as functions of $\boldsymbol{\psi}$:

$$\boldsymbol{\eta}_t = \boldsymbol{g}(\boldsymbol{\psi}_t) = \boldsymbol{F}_t' \boldsymbol{\beta}_t \quad \text{and} \quad \boldsymbol{\beta}_t = \boldsymbol{G}_t \boldsymbol{\beta}_{t-1} + \boldsymbol{\omega}_t, \quad \boldsymbol{\omega}_t \sim [\mathbf{0}, \boldsymbol{W}_t], \quad (4)$$

with $\boldsymbol{g} : \mathbb{R} \times \mathbb{R}^+ \rightarrow \mathbb{R}^2$, $\boldsymbol{g}(\boldsymbol{\psi}_t) = (g_1(\mu_t), g_2(\phi_t))$, \boldsymbol{F}_t is a $p \times 2$ matrix, where $p = p_1 + p_2$, with $p_i = \dim \boldsymbol{\beta}_{ti}$ and $\boldsymbol{\beta}_{ti} = (\beta_{ti1}, \dots, \beta_{tip_i})'$, $i = 1, 2$, the latent variables vector related to μ_t and ϕ_t . Depending on the specification of \boldsymbol{F}_t a broad class of models can be entertained. If $\boldsymbol{F}_t = \text{diag}(F_{t1}, F_{t2})$, different time series components and covariates are used to describe the time evolution of μ_t and ϕ_t through the link functions. Of course, they can also share some common regressors. The state parameters' evolution is described by a partially specified distribution, with $\boldsymbol{\omega}_t \sim [\mathbf{0}, \boldsymbol{W}_t]$, where $[a, b]$ denotes a distribution specified just by its first and second moments. The state parameters' initial information, $\boldsymbol{\beta}_0|D_0 \sim [\boldsymbol{m}_0, \boldsymbol{C}_0]$, is also partially specified with prior moments \boldsymbol{m}_0 and \boldsymbol{C}_0 .

Therefore equations in (3) and (4), together with the state parameters' initial information, define a class of partially specified models, where only the first and second prior moments are defined for the vector of latent components.

3 Inference in EDGLM

The class of models described by (3) and 4) extends the models treated in West et al. (1985) not only allowing the scale parameter to vary in time, but also modelling it through an additional link function. This extension implies that the original algorithm is not immediately applicable.

The conjugate updating algorithm of West et al. (1985) is extended, making estimation in this class of models feasible. The estimation is still based in the conjugate distribution and linear Bayes estimation, updating sequentially the state vector distributions at each time t , as in the original algorithm. At the end of this process, we obtain both the first and second posterior moments of latent states vectors and the posterior distribution of $(\psi_t|D_t)$ for each instant t .

In the next subsections, we review the main steps involved in the conjugate updating algorithm mainly to set up the notation, and propose a strategy to reduce the system dimension. We also discuss the forecasting distribution and conclude with some examples.

3.1 Extended Conjugate Updating

The conjugate updating algorithm is based on the steps: evolution, moments equating and updating. The evolution step involves obtaining the first and second moments of the state vectors prior distribution, $\mathbf{a}_t := E[\beta_t|D_{t-1}] = \mathbf{G}_t \mathbf{m}_{t-1}$ and $\mathbf{R}_t := Var(\beta_t|D_{t-1}) = \mathbf{G}_t \mathbf{C}_{t-1} \mathbf{G}_t' + \mathbf{W}_t$, given the posterior mean and variance at time $t-1$, \mathbf{m}_{t-1} , \mathbf{C}_{t-1} and the state evolution variance \mathbf{W}_t . The prior moments for the linear predictors follow immediately as: $\mathbf{f}_t := E[\eta_t|D_{t-1}] = \mathbf{F}_t' \mathbf{a}_t$ and $\mathbf{Q}_t := Var(\eta_t|D_{t-1}) = \mathbf{F}_t' \mathbf{R}_t \mathbf{F}_t$.

Alternatively, the prior moments of the linear predictor, $\eta_t = \mathbf{g}(\psi_t)$, can be obtained as functions of parameters defining the conjugate prior. Denote these prior moments as: $E[\eta_t|D_{t-1}] = \mathbf{h}(\boldsymbol{\tau}_t)$ and $Var(\eta_t|D_{t-1}) = \mathbf{H}(\boldsymbol{\tau}_t)$, where $\mathbf{h} : \mathbb{R}^3 \rightarrow \mathbb{R}^2$, $\mathbf{H} : \mathbb{R}^3 \rightarrow \mathcal{M}$ and \mathcal{M} is a set of symmetric positive definite 2×2 matrices and $\boldsymbol{\tau} = (\tau_0, \tau_1, \tau_2)'$ is the parameters vector of the conjugate prior.

We are facing a similar problem to the one posed by Poole and Raftery (2000) in the context of computer simulation models. There are two prior on the same quantity but based on different sources of information. This also occurs in the context of reaching consensus in the presence of multiple expert opinions. The analytic expressions of the above moments need to be equated to the linear predictors' numerical moments, previously obtained as functions of the prior moments of the states, providing the non-linear system of equations:

$$\mathbf{h}(\boldsymbol{\tau}_t) = \mathbf{f}_t \quad \text{and} \quad \text{vec}(\mathbf{H}(\boldsymbol{\tau}_t)) = \text{vec}(\mathbf{Q}_t). \quad (5)$$

where $\text{vec}(\mathcal{M})$ denotes the vectorization of the upper triangular matrix of a symmetric \mathcal{M} .

Note that the dimension of the involved vectors and matrices leads to a non-linear system with more equations than unknown quantities, so the system (5) does not provide a unique solution for the parameter

vector $\boldsymbol{\tau}_t$. Therefore it is necessary to introduce some criterion to reduce this large set of solutions to one compromise solution. A proposal to deal with this sort of dimension incompatibility in system (5) is treated in Section 3.2. This aims to answer the following query: What is the “best conjugate prior distribution” corresponding to the partially specified predictive distribution with mean \mathbf{f}_t and variance \mathbf{Q}_t ?

After observing a new datum, the prior parameters are straightforwardly updated. It follows from conjugacy that $\boldsymbol{\tau}_t$ can be updated according to expressions in (2), giving a new parameter vector $\boldsymbol{\tau}_t^*$. The linear predictors’ posterior moments can be obtained analogously to the system equations (5), given $\mathbf{f}_t^* = \mathbf{h}(\boldsymbol{\tau}_t^*)$ and $\mathbf{Q}_t^* = \mathbf{H}(\boldsymbol{\tau}_t^*)$, or analogously, $\text{vec}(\mathbf{Q}_t^*) = \text{vec}(\mathbf{H}(\boldsymbol{\tau}_t^*))$.

The observed information is propagated to the state vector using linear Bayes estimation (West and Harrison (1997), Chapter 4), since its distribution is only partially specified. Then, we obtain the posterior moments of $\boldsymbol{\beta}_t$, $\mathbf{m}_t = E[\boldsymbol{\beta}_t|D_t] = \mathbf{a}_t + \mathbf{R}_t \mathbf{F}_t \mathbf{Q}_t^{-1}(\mathbf{f}_t^* - \mathbf{f}_t)$ and $\mathbf{C}_t = \text{Var}(\boldsymbol{\beta}_t|D_t) = \mathbf{R}_t + \mathbf{R}_t \mathbf{F}_t \mathbf{Q}_t^{-1}[\mathbf{Q}_t^* - \mathbf{Q}_t] \mathbf{Q}_t^{-1} \mathbf{F}_t' \mathbf{R}_t$. The smoothed posterior moments of the latent states can be obtained in the same way as \mathbf{m}_t and \mathbf{C}_t , using linear Bayes estimation, as detailed in Souza (2013), resulting to expressions, $\mathbf{m}_t^s = E[\boldsymbol{\beta}_t|D_T] = \mathbf{m}_t + \mathbf{C}_t \mathbf{G}_{t+1}' \mathbf{R}_{t+1}^{-1}(\mathbf{m}_{t+1}^s - \mathbf{a}_{t+1})$ and $\mathbf{C}_t^s = \text{Var}(\boldsymbol{\beta}_t|D_T) = \mathbf{C}_t + \mathbf{C}_t \mathbf{G}_{t+1}' \mathbf{R}_{t+1}^{-1}(\mathbf{C}_{t+1}^s - \mathbf{R}_{t+1}) \mathbf{R}_{t+1}^{-1} \mathbf{G}_{t+1} \mathbf{C}_t$, where $\mathbf{m}_T^s := \mathbf{m}_T$ and $\mathbf{C}_T^s := \mathbf{C}_T$.

3.2 Dimensionality Reduction

To ensure the uniqueness of the vector $\boldsymbol{\tau}_t$ at each time considered in the algorithm, we need to reduce the dimensionality of the system (5). Several possibilities can be explored for this reduction, including arbitrary solutions such as ignoring some equations of the system (5). To avoid such arbitrariness we propose an alternative inspired on the generalized method of moments (Yin (2009)). Our main objective is to match the linear predictors’ moments and the conjugate prior moments preserving as much information provided by the system as possible. An optimum solution is obtained by minimizing the quadratic distance between the functional form that represents the difference between the numerical moments and the moment conditions described by its parameter vector, weighted by a weights matrix $\boldsymbol{\Omega}_k$ (where k is the dimension of the system) and zero. So, an optimum choice for the parameter vector $\boldsymbol{\tau}_t$ is the one that minimizes the function

$$\Delta_k(\boldsymbol{\tau}_t; \mathbf{f}_t, \mathbf{Q}_t)' \boldsymbol{\Omega}_k \Delta_k(\boldsymbol{\tau}_t; \mathbf{f}_t, \mathbf{Q}_t), \quad (6)$$

where $\Delta_k(\boldsymbol{\tau}_t; \mathbf{f}_t, \mathbf{Q}_t) = (\mathbf{f}_t - \mathbf{h}(\boldsymbol{\tau}_t), \text{vec}(\mathbf{Q}_t) - \text{vec}(\mathbf{H}(\boldsymbol{\tau}_t)))$ is a vectorial function and $\boldsymbol{\Omega}_k$ a positive definite weight matrix that specifies the importance of each equation condition in the estimation process.

Actually, since the weight matrix $\boldsymbol{\Omega}_k$ determines how each condition is weighted in the system solution, a simple choice is to take $\boldsymbol{\Omega}_k = \mathbf{I}_k$ (identity matrix of dimension k), which corresponds to considering all

the equations in system (5) on equal footing. Of course other choices for the matrix $\mathbf{\Omega}_k$ can be considered. Intuitively, the more accurate equations should be weighted more than the less accurate ones. A two-stage iterative procedure, described in Yin (2009), can be implemented to determine the “optimal” $\mathbf{\Omega}_k$ taking into account the observed data.

In summary, the proposed procedure can be implemented following the algorithm below:

Extended Conjugate Updating Algorithm:

At each time t

Step 1. evolution: given \mathbf{m}_{t-1} and \mathbf{C}_{t-1} ,

$$\mathbf{a}_t = \mathbf{G}_t \mathbf{m}_{t-1} \quad \text{and} \quad \mathbf{R}_t \mathbf{G}_t \mathbf{C}_{t-1} \mathbf{G}_t' + \mathbf{W}_t$$

$$\mathbf{f}_t = \mathbf{F}_t' \mathbf{a}_t \quad \text{and} \quad \mathbf{Q}_t = \mathbf{F}_t' \mathbf{R}_t \mathbf{F}_t.$$

Step 2. prior moment equating: obtain the prior parameter vector $\boldsymbol{\tau}_t$, solution of

$$\arg \min_{\boldsymbol{\tau}_t} \{ \Delta_k(\boldsymbol{\tau}_t; \mathbf{f}_t, \mathbf{Q}_t)' \mathbf{\Omega}_k \Delta_k(\boldsymbol{\tau}_t; \mathbf{f}_t, \mathbf{Q}_t) \}.$$

Step 3. posterior moments updating and equating: obtain $\boldsymbol{\tau}_t^*$ using equation (2) and calculate \mathbf{f}_t^* and \mathbf{Q}_t^* using $\boldsymbol{\tau}_t^*$ in equations (5).

Step 4. state updating: obtain $(\mathbf{m}_t, \mathbf{C}_t)$ via Linear Bayes estimation taking

$$\mathbf{m}_t = \mathbf{a}_t + \mathbf{R}_t \mathbf{F}_t \mathbf{Q}_t^{-1} (\mathbf{f}_t^* - \mathbf{f}_t) \quad \text{and}$$

$$\mathbf{C}_t = \mathbf{R}_t + \mathbf{R}_t \mathbf{F}_t \mathbf{Q}_t^{-1} [\mathbf{Q}_t^* - \mathbf{Q}_t] \mathbf{Q}_t^{-1} \mathbf{F}_t' \mathbf{R}_t.$$

3.3 Some Illustrative Examples

In this section we present examples involving the normal, the inverse Gaussian and the gamma distribution, leaving the discussion of the beta model to the next section. Our aim is to show the main functions involved in the Ef definition and their constraints.

3.3.1 Normal distribution with unknown mean and precision

Consider model (3), where $p(y_t | \mu_t, \phi_t)$ represents the density function of normal distribution with mean $\mu_t = \theta_t$ and variance ϕ_t^{-1} . In this case, $d_1(y_t) = y_t$, $d_2(y_t) = -\frac{y_t^2}{2}$ and $\rho(\theta_t, \phi_t) = \frac{1}{2} (\mu_t^2 \phi_t - \log(\phi_t))$, and the conjugate prior distribution takes the form

$$p(\mu_t, \phi_t | D_{t-1}) \propto \exp \{ \phi_t [\mu_t \tau_{1t} + \tau_{2t}] - \tau_{0t} \rho(\mu_t, \phi_t) \}, \quad \mu_t \in \mathbb{R}, \quad \phi_t \in \mathbb{R}^+$$

which represents the kernel of the density function of the normal-gamma distribution with parameters $\frac{\tau_{1t}}{\tau_{0t}}$, τ_{0t} , $\frac{\tau_{0t}+1}{2}$ and $-\frac{\tau_{1t}^2}{2\tau_{0t}} - \tau_{2t}$.

Using the natural link functions $\eta_{1t} = g_1(\mu_t) = \mu_t$ and $\eta_{2t} = g_2(\phi_t) = \log(\phi_t)$, and the crude approximation of the digamma function, $\psi(x) = \log(x) + O(x)$, $x > x_0$, to evaluate the moments of the linear predictor η_{2t} , it follows that moment conditions are represented as in the functional form

$$\begin{aligned} \Delta_k(\boldsymbol{\tau}_t; \mathbf{f}_t, \mathbf{Q}_t) = & \left(f_{1t} - \frac{\tau_{1t}}{\tau_{0t}}, f_{2t} - \log \left(\frac{\tau_{0t}+1}{2} \left[-\frac{\tau_{1t}^2}{2\tau_{0t}} - \tau_{2t} \right]^{-1} \right), \right. \\ & \left. q_{11t} - \left[-\frac{\tau_{1t}^2}{2\tau_{0t}} - \tau_{2t} \right] \tau_{0t}^{-1} \left[\frac{\tau_{0t}+1}{2} - 1 \right]^{-1}, q_{22t} - \frac{2}{\tau_{0t}+1} \right). \end{aligned} \quad (7)$$

Therefore $\boldsymbol{\tau}_t$ is obtained as the solution that minimizes the associated quadratic form.

Note that in this example the prior covariance of the linear predictors (η_{1t}, η_{2t}) , at each time t , are zero, which indicates that \mathbf{Q}_t is a diagonal matrix. In fact, it means that η_1 is orthogonal to η_2 given D_{t-1} , so the system reduces to four equations. Nevertheless solving system (7) is not a trivial minimization problem since we need to ensure that all involved moments are well defined, in the sense that at each algorithm's iteration, τ_{0t} , τ_{1t} and τ_{2t} generate non-negative variances. In this particular example, the minimization with respect to the vector $\boldsymbol{\tau}_t$ must satisfy the restrictions $\tau_{0t} > 1$ and $\tau_{2t} < -\frac{\tau_{1t}^2}{2\tau_{0t}}$, assuming that the first and second moments of expression (7) are well defined.

3.3.2 Inverse Gaussian distribution

Suppose that $p(y_t|\mu_t, \phi_t)$ represents the density function of inverse normal distribution with mean μ_t and variance $\frac{\mu_t^3}{\phi_t}$ in model (3). It is very easy to show that this model is a member of the exponential family, taking $d_1(y_t) = -y_t$, $d_2(y_t) = -\frac{1}{2y_t}$, $\rho(\mu_t, \phi_t) = -\left[\frac{\phi_t}{\mu_t} + \frac{1}{2} \log(\phi_t) \right]$ and $a(y) = (2\pi y_t^3)^{-1/2}$. In this case, the conjugate prior distribution for the observational model is

$$p(\mu_t, \phi_t | D_{t-1}) \propto \exp \left\{ -\phi_t \left[\frac{1}{2\mu_t^2} \tau_{1t} + \frac{1}{2} \tau_{2t} \right] + \tau_{0t} \rho(\mu_t, \phi_t) \right\}, \quad \mu_t > 0, \phi_t > 0. \quad (8)$$

As explained in Banerjee and Bhattacharyya (1979), conditional to ϕ_t , μ_t^{-1} follows a normal distribution truncated at zero; and, conditional to μ_t , ϕ_t follows a gamma distribution. On the other hand, $p(\mu_t, \phi_t | D_{t-1})$ does not have an analytically known form, as far as we know, so we approximate its mean and variance by the mode $(\tilde{\mu}_t, \tilde{\phi}_t)'$ and the inverse curvature matrix \tilde{V}_t of the conjugate prior distribution (8) evaluated at the mode point, respectively, getting

$$\begin{pmatrix} \tilde{\mu}_t \\ \tilde{\phi}_t \end{pmatrix} = \begin{pmatrix} \tau_{1t}/\tau_{0t} \\ \tau_{0t} \left[\tau_{2t} - \frac{\tau_{0t}^2}{\tau_{1t}} \right]^{-1} \end{pmatrix} \quad \text{and} \quad \tilde{V}_t = \begin{bmatrix} \frac{\tilde{\mu}_t^3}{\tau_{0t}\tilde{\phi}_t} & 0 \\ 0 & \frac{2\tilde{\phi}_t^2}{\tau_{0t}} \end{bmatrix}.$$

Using the link functions $g_1(\mu_t) = \log(\mu_t)$ and $g_2(\phi_t) = \log(\phi_t)$, and taking first-order Taylor approximations of these functions around $(\tilde{\mu}_t, \tilde{\phi}_t)'$, we obtain the mode and curvature of the linear predictors.

Then to equate the numerical moments of the linear predictors with those obtained using their conjugate prior we must solve the system of equations

$$\begin{aligned} \Delta_k(\boldsymbol{\tau}_t; \mathbf{f}_t, \mathbf{Q}_t) = & \left(f_{1t} - \log(\tau_{1t}) + \log(\tau_{0t}), f_{2t} - \log\left(\frac{\tau_{0t}\tau_{1t}}{\tau_{1t}\tau_{2t} - \tau_{0t}^2}\right), \right. \\ & \left. q_{11t} - \frac{\tau_{1t}\tau_{2t}}{\tau_{0t}^3} + \frac{1}{\tau_{0t}}, q_{22t} - \frac{2}{\tau_{0t}} \right). \end{aligned} \quad (9)$$

The optimization problem (7), based on $\boldsymbol{\Delta}_k$ like in (9), must satisfy the constraints $\tau_{2t} > \frac{\tau_{0t}}{\tau_{1t}}$ with $\tau_{0t}, \tau_{1t} > 0$ in order to ensure that all variances are positive.

3.3.3 Gamma distribution

Let $y_t|\mu_t, \phi_t$ denote the density function of the gamma distribution, with mean μ_t and variance $\frac{\mu_t^2}{\phi_t}$. The quantities defining this member of the two-parameter exponential family are: $\theta_t = \frac{1}{\mu_t}$, $d_1(y_t) = -y_t$, $d_2(y) = \log(y_t)$ and $\rho(\theta_t, \phi_t) = \log(\Gamma(\phi_t)) - \phi_t \log\left(\frac{\phi_t}{\mu_t}\right)$ and, therefore, its conjugate prior distribution is given by

$$p(\mu_t, \phi_t | D_{t-1}) \propto \exp\left\{\phi_t \left[-\frac{1}{\mu_t}\tau_{1t} + \tau_{2t}\right] - \tau_{0t}\rho(\theta_t, \phi_t)\right\}, \quad \mu_t > 0, \phi_t > 0. \quad (10)$$

Since the prior distribution does not represent a known distribution, as far as we know, we opt to use its mode and the inverse curvature matrix of the conjugate prior distribution (10) in place of its mean and variance. Using the logarithmic link functions for both parameters, we get

$$\begin{aligned} E(\eta_{1t} | D_{t-1}) &\approx \log(\tau_{1t}) - \log(\tau_{0t}) \quad \text{and} \quad E(\eta_{2t} | D_{t-1}) \approx \log\left(\frac{\tau_{0t}}{2\left[\tau_{0t}\log\left(\frac{\tau_{1t}}{\tau_{0t}}\right) - \tau_{2t}\right]}\right) \\ \text{Var}(\eta_{1t} | D_{t-1}) &\approx \frac{2}{\tau_{0t}^2} \left[\tau_{0t}\log\left(\frac{\tau_{1t}}{\tau_{0t}}\right) - \tau_{2t}\right] \quad \text{and} \quad \text{Var}(\eta_{2t} | D_{t-1}) \approx \frac{2}{\tau_{0t}}. \end{aligned} \quad (11)$$

Moreover, taking a first order Taylor approximation of the function $\mathbf{g}(\mu_t, \phi_t) = (\log(\mu_t)\log(\phi_t))$ around the mode of (10), we obtain the covariance of the linear predictors as

$$\text{Cov}(\eta_{1t}, \eta_{2t} | D_{t-1}) \approx \log(\tilde{\mu}_t)\log(\tilde{\phi}_t) - [\log(\tilde{\mu}_t)][\log(\tilde{\phi}_t)] = 0.$$

Comparing the numerical moments obtained for linear predictors through the dynamic model with those obtained by conjugation (expressions (11)), we obtain the functional form

$$\Delta_k(\boldsymbol{\tau}_t; \mathbf{f}_t, \mathbf{Q}_t) = \left(f_{1t} - \log\left(\frac{\tau_{1t}}{\tau_{0t}}\right), f_{2t} - \log\left(\frac{\tau_{0t}}{2\left[\tau_{0t}\log\left(\frac{\tau_{1t}}{\tau_{0t}}\right) - \tau_{2t}\right]}\right), q_{11t} - \frac{1}{\tau_{0t}\tilde{\phi}_t}, q_{22t} - \frac{2}{\tau_{0t}} \right), \quad (12)$$

whose quadratic distance with respect to zero (possibly weighted by a weights matrix $\mathbf{\Omega}_k$) can be minimized by imposing the constraints $\tau_{0t} > 0$, $\tau_{1t} > 0$ and $\frac{\tau_{2t}}{\tau_{0t}} > f_{1t}$, which ensures that the moments up to second order associated with the conjugate prior distribution (10) are well defined.

3.4 Forecasting

Assume that our interest is to forecast some future observation, for example, at instant $t + h$ (for some integer h), based on all observations until instant t . Making use of exponential family's proprieties, it follows from conjugacy that

$$p(y_{t+h}|D_t) = a(y_{t+h}) \frac{\kappa(\boldsymbol{\tau}_{t+h})}{\kappa(\boldsymbol{\tau}_{t+h}^*)}, \quad (13)$$

where $\kappa(\boldsymbol{\tau}_{t+h})$ and $\kappa(\boldsymbol{\tau}_{t+h}^*)$ are the normalization constants involved in the definition of the prior and the posterior distribution of the vector $(\theta_{t+h}, \phi_{t+h})$, respectively. Here, the parameter vector $\boldsymbol{\tau}_{t+h} = (\tau_{0,t+h}, \tau_{1,t+h}, \tau_{2,t+h})$ can be obtained analogously to that discussed in Section 3.2, by solving the optimization problem

$$\arg \min_{\boldsymbol{\tau}_{t+h}} \{ \Delta_k(\boldsymbol{\tau}_{t+h}; \mathbf{f}_t(h), \mathbf{Q}_t(h))' \mathbf{\Omega}_k \Delta_k(\boldsymbol{\tau}_{t+h}; \mathbf{f}_t(h), \mathbf{Q}_t(h)) \}, \quad (14)$$

given the recursive relation between the linear predictor moments

$$\begin{aligned} \boldsymbol{\eta}_t | D_t &\sim [\mathbf{f}_t(h), \mathbf{Q}_t(h)], & \mathbf{f}_t(h) &= \mathbf{F}'_{t+h} \mathbf{a}_t(h), \\ & & \mathbf{Q}_t(h) &= \mathbf{F}'_{t+h} \mathbf{R}_t(h) \mathbf{F}_{t+h}, \quad \text{with} \end{aligned}$$

$$\begin{aligned} \mathbf{a}_t(h) &= \mathbf{G}_{t+h} \mathbf{a}_t(h-1), & \mathbf{R}_t(h) &= \mathbf{G}_{t+h} \mathbf{R}_t(h-1) \mathbf{G}'_{t+h} + \mathbf{W}_{t+h} & \text{and} \\ \mathbf{a}_t(0) &= \mathbf{m}_t, & \mathbf{R}_t(0) &= \mathbf{C}_t. \end{aligned}$$

Note that the vector $\boldsymbol{\tau}_{t+h}^* = (\tau_{0,t+h}^*, \tau_{1,t+h}^*, \tau_{2,t+h}^*)$ is directly obtained like in the relations represented in (2).

In cases in which the constants $\kappa(\cdot)$ do not have known analytical form, we must use some numerical integration method to approximate them. In this work, Laplace approximations are used to solve such integrals. All methods are implemented with the aid of routines available in the free software R (Team (2011)), like the optimization function `nlminb` and the function `fdHess` which numerically approximate gradient and Hessian functions. Furthermore, to improve the quality of the approaches, we use a new parameterization for the involved prior distributions in terms of their linear predictors η_{1t} and η_{2t} , integrating new parameters along the real line. See the next section for an example.

4 Case Studies

In this section, two applications are presented to illustrate the performance of the proposed method. In both cases we suppose that the observations follow a beta distribution. The first one models unemployment rates in Brazil, using a data set that presents a trend component and a stable seasonal pattern. Our main interest in this first application is to illustrate the importance of dynamically modelling the precision parameter. The second one considers some macroeconomic variables of the United Kingdom, viewed as compositional data. In this example our aim is to show the importance of modelling the data in their original scale. To start this section, we show the main developments concerning the beta model used in both applications. The implementations are carried out through the R software and more details is discussed below.

4.1 Dynamic beta model components

Consider now $p(y_t|\mu_t, \phi_t)$ as the density function of a beta distribution in model (3), parameterized in terms of its mean μ_t and its variance $\frac{\mu(1-\mu)}{\phi}$. In this case, using conjugacy in the exponential family,

$$p(\mu_t, \phi_t|D_{t-1}) \propto \exp\{\phi_t[\mu_t\tau_{1t} + \tau_{2t}] + \tau_{0t}\rho(\mu_t, \phi_t)\}, \quad 0 < \mu_t < 1, \quad \phi_t > 0. \quad (15)$$

where $\rho(\mu_t, \phi_t) = -\log\left(\frac{\Gamma(\phi_t)}{\Gamma(\phi_t\mu_t)\Gamma(\phi_t(1-\mu_t))}\right)$. Taking $g_1(\mu_t) = \text{logit}(\mu_t)$ and $g_2(\phi_t) = \log(\phi_t)$ and approximating first and second moments of (15), respectively, by the mode $(\tilde{\mu}_t, \tilde{\phi}_t)'$ and the inverse curvature matrix evaluated at the mode, we get

$$\begin{aligned} E(\eta_{1t}|D_{t-1}) &\approx \text{logit}(\tilde{\mu}_t) = \frac{\tau_{1t}}{\tau_{0t}} \quad \text{and} \quad \text{Var}(\eta_{1t}|D_{t-1}) \approx \frac{1}{\tau_{0t}\tilde{\mu}_t(1-\tilde{\mu}_t)\tilde{\phi}_t} \\ E[\eta_{2t}|D_{t-1}] &\approx \log(\tilde{\phi}_t) = \log\left(\frac{\tau_{0t}}{2\{\tau_{0t}\log(1-\tilde{\mu}_t) - \tau_{2t}\}}\right) \quad \text{and} \quad \text{Var}(\eta_{2t}|D_{t-1}) \approx \frac{2}{\tau_{0t}}. \end{aligned} \quad (16)$$

The functional form (7) to be minimized depends on the vector function

$$\Delta_k(\boldsymbol{\tau}_t; \mathbf{f}_t, \mathbf{Q}_t) = \left(f_{1t} - \frac{\tau_{1t}}{\tau_{0t}}, f_{2t} - \log\left(\frac{\tau_{0t}}{2\{\tau_{0t}\log(1-\tilde{\mu}_t) - \tau_{2t}\}}\right), q_{11t} - \frac{1}{\tau_{0t}\tilde{\mu}_t(1-\tilde{\mu}_t)\tilde{\phi}_t}, q_{22t} - \frac{2}{\tau_{0t}} \right), \quad (17)$$

whose minimum must be obtained by imposing restrictions $\tau_{0t} > 0$ and $\tau_{2t} > -\tau_{0t}\log(1 + \exp\{f_{1t}\})$, since we are imposing the condition that $\text{Cov}(\eta_{1t}, \eta_{2t}|D_{t-1}) = 0$, such as in the gamma case.

It is worth noting that although the beta distribution has a conjugated prior represented in equation (15), it does not have a known analytical form, as far as we know. So, to find its normalization constant we need to approximate the integral

$$\kappa(\tau_{0t}, \tau_{1t}, \tau_{2t})^{-1} = \int_0^\infty \int_0^1 \exp\{\phi_t[\mu_t\tau_{1t} + \tau_{2t}] + \tau_{0t}\rho(\mu_t, \phi_t)\} d\mu_t d\phi_t,$$

by using a Laplace approximation for its expression. In fact, by changing the variables of the integral in (18) in terms of η_{1t} and η_{2t} , we can approximate it as

$$\begin{aligned}\kappa(\tau_{0t}, \tau_{1t}, \tau_{2t})^{-1} &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \exp \left\{ e^{\eta_{2t}} \left[\frac{e^{\eta_{1t}}}{1 + e^{\eta_{1t}}} \tau_{1t} + \tau_{2t} \right] - \tau_{0t} \rho'(\eta_{1t}, \eta_{2t}) \right\} d\eta_{1t} d\eta_{2t} \\ &\approx \sqrt{2\pi} |\tilde{V}_t|^{\frac{1}{2}} \exp\{L_t(\tilde{\eta}_{1t}, \tilde{\eta}_{2t})\},\end{aligned}$$

where $\rho'(\eta_{1t}, \eta_{2t}) = \log(\Gamma(e^{\eta_{2t}})) - \log\left(\Gamma\left(e^{\eta_{2t}} \frac{e^{\eta_{1t}}}{1 + e^{\eta_{1t}}}\right) \Gamma\left(e^{\eta_{2t}} \frac{1}{1 + e^{\eta_{1t}}}\right)\right)$, $L_t(\eta_{1t}, \eta_{2t}) = e^{\eta_{2t}} \left[\frac{e^{\eta_{1t}}}{1 + e^{\eta_{1t}}} \tau_{1t} + \tau_{2t} \right] + \tau_{0t} \rho'(\eta_{1t}, \eta_{2t})$ and $\tilde{V}_t = -[\nabla^2 L_t(\eta_{1t}, \eta_{2t})]^{-1} \Big|_{(\eta_{1t}, \eta_{2t}) = (\tilde{\eta}_{1t}, \tilde{\eta}_{2t})}$ is the Hessian matrix of $L_t(\eta_{1t}, \eta_{2t})$, applied in its mode $(\tilde{\eta}_{1t}, \tilde{\eta}_{2t})$.

Using the R software, the mode $(\tilde{\eta}_{1t}, \tilde{\eta}_{2t})$ and the Hessian matrix \tilde{V}_t can be easily obtained using, respectively, the functions `nlminb` and `fdHess`, using the expression $L_t(\eta_{1t}, \eta_{2t})$ as the argument.

4.2 Unemployment rates in Brazil

The data for this example was collected by the Brazilian Institute of Geography and Statistics (IBGE: <http://www.ibge.gov.br/>) through its Monthly Employment Survey and deals with monthly unemployment rates of working-age people in the major metropolitan regions of Brazil, namely the metropolitan areas of Recife, Salvador, Belo Horizonte, Rio de Janeiro, São Paulo and Porto Alegre. The monthly unemployment rates of working-age people from March 2002 to December 2011, in a total of 118 observations, can be seen in Figure 1. This time series clearly exhibits components of trend and seasonality.

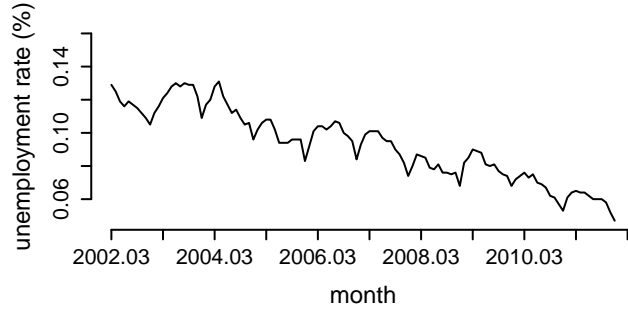


Figure 1: Unemployment rates of working-age people in the major metropolitan regions of Brazil from March 2002 to December 2011.

It is well known that the yearly seasonal behaviour in this time series is mainly due to temporary jobs created by holiday seasons and school vacations, as mentioned by da Silva et al. (2011). Considering these

factors, we analysed the data set through a dynamic beta model, where the observational mean evolve as a second-order polynomial model with seasonal effect. Unlike da Silva et al. (2011), we assume a more parsimonious model, where seasonality is represented by a one-harmonic model and we assume that the precisions can evolve dynamically in time. Additionally, we assume that the latent variables associated with means and precisions evolve in time independently, taking the matrices \mathbf{F}'_t , \mathbf{G}_t , \mathbf{W}_t and \mathbf{C}_0 as block diagonal matrices of the form $\mathbf{F}'_t = \text{diag}(\mathbf{F}'_{1t}, \mathbf{F}'_{2t})$, $\mathbf{G}_t = \text{diag}(\mathbf{G}_{1t}, \mathbf{G}_{2t})$, $\mathbf{W}_t = \text{diag}(\mathbf{W}_{1t}, \mathbf{W}_{2t})$ and $\mathbf{C}_0 = \text{diag}(\mathbf{C}_{10}, \mathbf{C}_{20})$, where the matrices related to the dynamics of the observational means are given by

$$\mathbf{F}_{1t} = (1, 0, 1, 0)' \quad \text{and} \quad \mathbf{G}_{1t} = \begin{pmatrix} \mathbf{J}_2(1) & \mathbf{0} \\ \mathbf{0} & \mathbf{J}_2(1, \omega) \end{pmatrix}, \quad \forall t, \quad \text{where}$$

$$\mathbf{J}_2(1) = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{J}_2(1, \omega) = \begin{pmatrix} \cos(\omega) & \sin(\omega) \\ -\sin(\omega) & \cos(\omega) \end{pmatrix}, \quad \omega = \frac{2\pi}{12}.$$

To model the dispersions, we assume a first order dynamic model, taking $\mathbf{F}_{2t} = \mathbf{G}_{2t} = \mathbf{I}$, $\forall t$, in order to allow precision parameter ϕ_t to vary in time through the introduction of a random error.

We chose to specify the error evolution covariance matrices, \mathbf{W}_t , $t \in \{1, \dots, T\}$, through the use of multiple discount factors assuming \mathbf{W}_t to be a block diagonal matrix whose blocks are associated with mean level and trend and seasonal components, and a precision level component, taking $\mathbf{D} = \text{blockdiag}\{\delta_{\mu,lt}^{-1/2} \mathbf{I}_2, \delta_{\mu,s}^{-1/2} \mathbf{I}_2, \delta_{\phi,l}^{-1/2}\}$, where $\delta_{\mu,lt}$, $\delta_{\mu,s}$ and $\delta_{\phi,l}$ are discount factors associated with the respective blocks of components by replacing the expression of \mathbf{R}_t in the evolution step of the algorithm with the form $\mathbf{R}_t = \mathbf{D} \mathbf{G}_t \mathbf{C}_{t-1} \mathbf{G}'_t \mathbf{D}$.

Different combinations of discount factors were tried and we selected the one that provided the best performance according to some alternative model selection criteria like the mean squared error (MSE) based on one-step-ahead forecasting, the joint log-likelihood (LL) and the log-observed predictive density (LPD), excluding the first 18 observations, taken as a learning period. Using the selected discount factors, namely, $\delta_{\mu,lt} = 0.90$, $\delta_{\mu,s} = 0.95$ and $\delta_{\phi,l} = 0.90$, we obtained the model parameter estimates and the one-step-ahead predictive distributions for the unemployment rates during the period from September 2003 to December 2011 at each instant, using expression (13) as discussed in the previous subsection with the aid of the R routines `nlminb` and `fdHess`.

In Figure 2, it is possible to observe the filtered ($E[\beta_t | D_{t-1}]$) and the smoothed estimated state variable means ($E[\beta_t | D_T]$) related to the observational mean components, describing level, trend and seasonality, respectively; and the state variable associated with the observational precision. In fact, there is a clearly decreasing trend in the data as well as a seasonal behaviour like observed in Figure 1.

Regarding the precision structure, the small growth of the state variable β_{5t} over time can indicate that as new information is incorporated in the estimation process, the accuracy of the model increases.

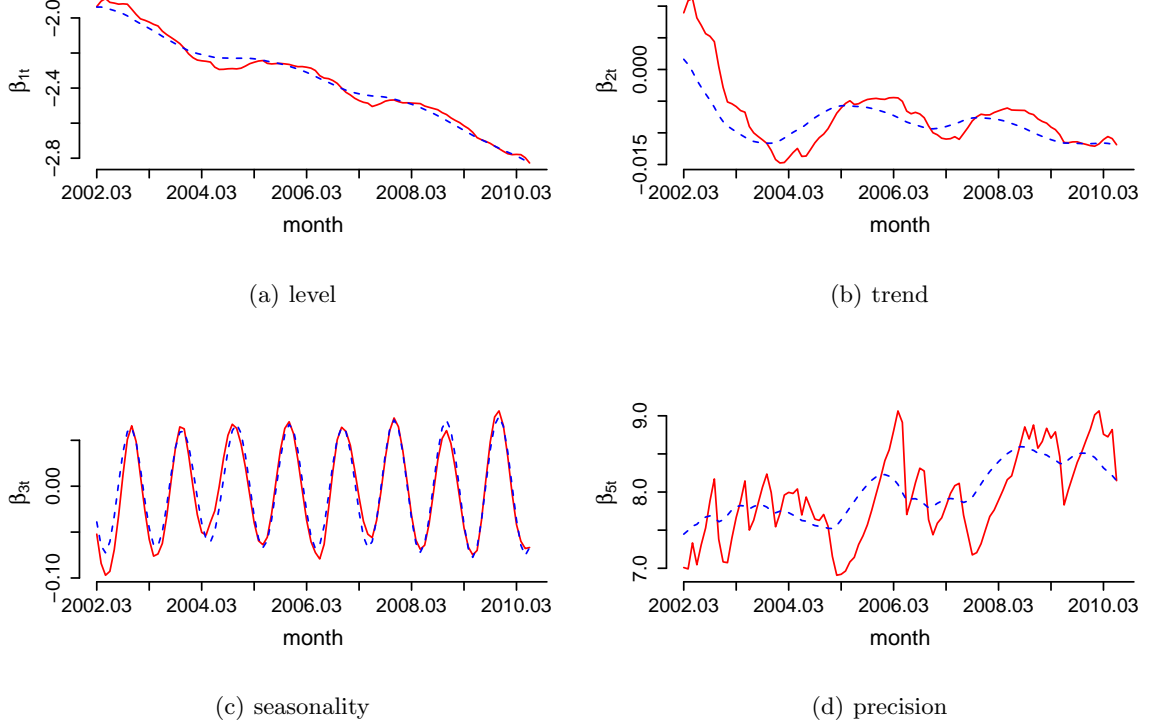


Figure 2: Filtered (solid line) and smoothed (dashed line) latent states means for application using unemployment data.

It can be seen in Figures 3 and 4 that the method generated satisfactory results, since both estimated means (the filtered ones $E[\mu_t|D_t]$,) and one-step-ahead predictive distribution means ($E[y_t|D_{t-1}]$) follow the behavior of the real data series, as illustrated by Figures 3 and 4, respectively. Also note that the estimated 95% HPD credibility intervals for the one-step-ahead predictive distributions, represented by the dashed red lines in Figure 4, are well concentrated and contain the true value of the observations in all considered instances. The point and interval estimates for the predictive distributions considered in the last six instants can be seen in Table 1.

To illustrate the importance of dynamic modelling for the precision parameter model, we completed this application by comparing its results with those obtained using a similar model in which we assumed that $\phi_t = \phi, \forall t$, taking null precision evolution errors in matrix \mathbf{W}_t . Figure 4 compares the interval estimates for the one-step-ahead predictive distribution obtained considering both models. Note that

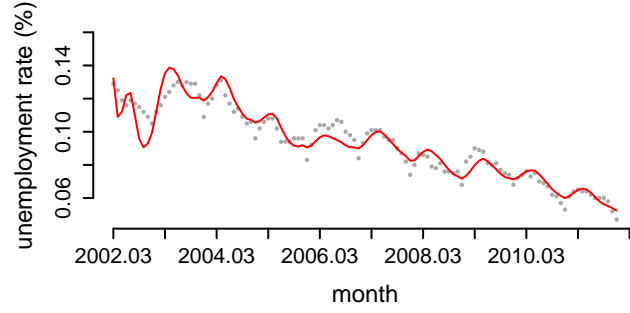


Figure 3: Filtered estimated observational means (solid line) based on latent states posterior means obtained for application using unemployment data. The gray points represent the true observations.

Month	y_t	Mean	Mode	IC _{95%}
2011.07	0.060	0.062	0.061	[0.054 , 0.070]
2011.08	0.060	0.059	0.058	[0.052 , 0.066]
2011.09	0.060	0.056	0.056	[0.050 , 0.063]
2011.10	0.058	0.055	0.055	[0.048 , 0.061]
2011.11	0.052	0.054	0.055	[0.048 , 0.061]
2011.12	0.047	0.054	0.054	[0.048 , 0.060]

Table 1: Point and interval estimates of unemployment rates for the period from July 2011 to December 2011, based on the predictive distributions $p(y_t|D_{t-1})$.

intervals based on a model with ϕ fixed in time (represented by the shaded area in the graph) are less concentrated, indicating that there was a gain with respect to accuracy of the predictive distributions in this case, in which we considered the dynamic modelling of the precision structure.

4.3 Expenditure shares in the U.K. economy

As a second illustration of the proposed methodology, we apply the new method to a real data set concerning expenditures in the UK economy for the period 1955 to 2012. The quarterly data, obtained from the U.K. Office of National Statistics web page (<http://www.statistics.gov.uk/>), deal with the costs of the economy, whose composition is described by consumption (c), investment (i), government expenditure (g) and export (e) shares of U.K. gross final expenditure.

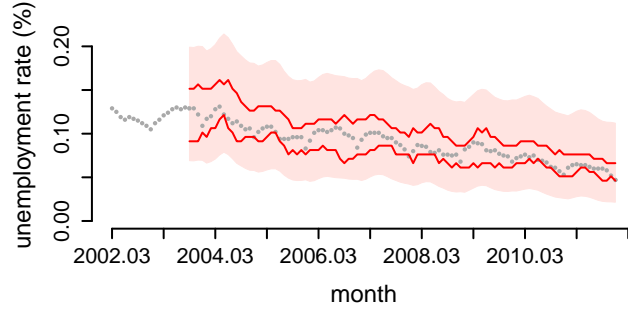


Figure 4: One-step-ahead predictive intervals from September 2003 to December 2011, based on the predictive distributions $p(y_t|D_{t-1})$ of alternative models: the shaded region represents the 95% HPD predictive credibility intervals based on a model with ϕ fixed in time, whereas the solid lines represent the predictive distribution credibility intervals based on a model with a dynamic precision parameter structure. The gray points are the observed unemployment rates.

Despite the compositional nature of the data, in order to use the class of models discussed in this article, which includes only univariate observational distributions in the exponential family, we analysed each of the rate series separately through a generalized dynamic model whose observations follow the beta distributions, and for which we assumed different mean and precision structures. We denote the proposed models by the mnemonic $\text{Var}(l)$ and $\text{Pol}(l)$, meaning a vectorial autorregressive component and a polynomial trend, respectively, where l is the order of the correspondent model. This models were combined to model the transformed observational mean and the transformed observational precision in different forms. For each case, as in the previous application, we assumed that the latent variables associated with means and precisions model evolve in time independently, taking the matrices \mathbf{F}'_t , \mathbf{G}_t , \mathbf{W}_t and \mathbf{C}_0 as block diagonal matrices. Under this hypothesis, three different structures were considered for the class of models represented by (3) and (4):

- **Var(2)Pol(0)** - Second order VAR model for the transformed observational mean and constant for the transformed observational precision:

For the means structure, we assumed that each series can be explained by all the other series, taking

two lags in time, assuming a second-order VAR model. For the precision structure, we assumed that each series has a constant accuracy in time, taking

$$\mathbf{F}_{1t} = \begin{pmatrix} 1, & x_{\mathbf{c},t-1}, & x_{\mathbf{g},t-1}, & x_{\mathbf{i},t-1}, & x_{\mathbf{e},t-1}, & x_{\mathbf{c},t-2}, & x_{\mathbf{g},t-2}, & x_{\mathbf{i},t-2}, & x_{\mathbf{e},t-2} \end{pmatrix}',$$

$$\mathbf{G}_{1t} = \mathbf{I}_9, \quad \mathbf{W}_{1t} = \mathbf{0}, \quad \text{and} \quad \mathbf{F}_{2t} = \mathbf{G}_{2t} = 1, \quad \mathbf{W}_{2t} = 0, \quad \forall t \in \{1, \dots, T\},$$

where $x_{\mathbf{c},\cdot}$, $x_{\mathbf{g},\cdot}$, $x_{\mathbf{i},\cdot}$, $x_{\mathbf{e},\cdot}$, represent, respectively, the rates of consumption, government expenditure, investment and exports in previous instants.

- **Var(2)Pol(1)** - Transformed observational mean modelled by a second order VAR and precision with a first order dynamic structure:

As in the previous case, we assumed means explained by a second-order VAR model, but in this case we allowed the precisions to vary in time according to a first order polynomial model taking

$$\mathbf{F}_{1t} = \begin{pmatrix} 1, & x_{\mathbf{c},t-1}, & x_{\mathbf{g},t-1}, & x_{\mathbf{i},t-1}, & x_{\mathbf{e},t-1}, & x_{\mathbf{c},t-2}, & x_{\mathbf{g},t-2}, & x_{\mathbf{i},t-2}, & x_{\mathbf{e},t-2} \end{pmatrix}',$$

$$\mathbf{G}_{1t} = \mathbf{I}_9, \quad \mathbf{W}_{1t} = \mathbf{0}, \quad \text{and} \quad \mathbf{F}_{2t} = \mathbf{G}_{2t} = 1, \quad \mathbf{W}_{2t} \neq 0, \quad \forall t \in \{1, \dots, T\},$$

where, again, $x_{\mathbf{c},\cdot}$, $x_{\mathbf{g},\cdot}$, $x_{\mathbf{i},\cdot}$, $x_{\mathbf{e},\cdot}$, represent, respectively, the rates of consumption, government expenditure, investment and exports in previous instants.

- **Pol(2)Pol(1)** - Polynomial models for both mean and precision structures:

For the means we assumed a second-order model in which we considered level and trend for each of the series and a first-order structure for the precisions, taking

$$\mathbf{F}_{1t} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad \mathbf{G}_{1t} = \mathbf{J}_2(1) = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \quad \mathbf{W}_{1t} \neq \mathbf{0} \quad \text{and}$$

$$\mathbf{F}_{2t} = \mathbf{G}_{2t} = 1, \quad \mathbf{W}_{2t} \neq 0 \quad \forall t \in \{1, \dots, T\}.$$

As in the previous application, we chose to specify the covariance matrices through the use of multiple discount factors, assuming block diagonal matrices, whose blocks are associated with the respective components (level and trend in the case of second-order model and level in the order 1 model) in polynomial models. More specifically, considering, for example, the **Pol(2)Pol(1)** structure, we used a block diagonal discount matrix of the form $\mathbf{D} = \text{blockdiag}\{\delta_{\mu,lt}^{-1/2} \mathbf{I}_2, \delta_{\phi,l}^{-1/2}\}$, where $\delta_{\mu,lt}$ is the discount factor associated with mean level and trend components and $\delta_{\phi,l}$ is the discount factor associated with

precision level components, substituting the expression of \mathbf{R}_t in the evolution step of the algorithm for the form $\mathbf{R}_t = \mathbf{D}\mathbf{G}_t\mathbf{C}_{t-1}'\mathbf{G}_t'\mathbf{D}$, as discussed in Chapter 6 of West and Harrison (1997).

For each of the rate series and for each of the dynamic structures assumed, different combinations of discount factors values were used, so we selected the one that provided the best data fit according to the mean squared error (MSE) based on one-step-ahead forecasting, the joint log-likelihood (LL) and the log-observed predictive density (LPD) of each series, excluding the first 31 observations (taken as learning sample). For this application, different combinations of values 0.90, 0.95 and 0.98 were taken for the discount factors and, for all assumed dynamic structures, models with smaller values, namely $\delta_{\mu,lt} = \delta_{\phi,l} = 0.90$, outperformed. Table 2 reports adjustment measures for the different dynamic models. It can be seen that, according to the criteria used, the model that supposes a second-order Var structure for the mean and a first order structure for the precision performs better with lower MSE and values and higher LL and LPD values, which makes sense since the Var structure capturing the relationship between the different rate series and allows the precision model structure to vary in time, giving greater flexibility to the model.

consumption				investment		
Model	MSE	LL	LPD	MSE	LL	LPD
VAR(2)Pol(0)	0.400e-4	745.507	724.151	0.525e-4	714.757	694.131
VAR(2)Pol(1)	0.400e-4	759.776	725.504	0.505e-4	730.774	698.827
Pol(2)Pol(1)	0.838e-4	706.711	653.134	1.083e-4	691.014	636.617

government expenditure				export		
Model	MSE	LL	LPD	MSE	LL	LPD
Var(2)Pol(0)	0.105e-4	881.688	857.297	0.495e-4	726.010	702.890
Var(2)Pol(1)	0.101e-4	899.035	862.609	0.463e-4	745.200	710.554
Pol(2)Pol(1)	0.453e-4	761.649	707.335	1.127e-4	676.193	621.606

Table 2: MSE based on one-step-ahead forecasting, joint log-likelihood (LL) and log-observed predictive density (LPD) based on consumption, investment, government expenditure and export rates for the period from 1963.2 to 2012.3, obtained from different models.

Once the Var(2)Pol(1) model was selected, we estimated the parameters and the one-step-ahead predictive distributions for the four rate series during the period 1963.2 to 2012.3, as represented by Figures 5 and 6, respectively. In Figure 5 it is possible to observe the point estimates for the observational means of each series (the filtered ones $E[\mu_t|D_t]$). Note that for all analysed series, the estimated means closely parallel the behaviour of the data series. Similar behaviour can also be observed for the estimated predictive mean's ($E[y_t|D_{t-1}]$), shown in Figure 6. It can also be seen that the estimated 95% HPD credibility intervals for the one-step-ahead predictive distributions are well concentrated, containing the true observation values in most cases. Point and interval predictive estimates for investment rates for some considered instants can be seen in Table 3.

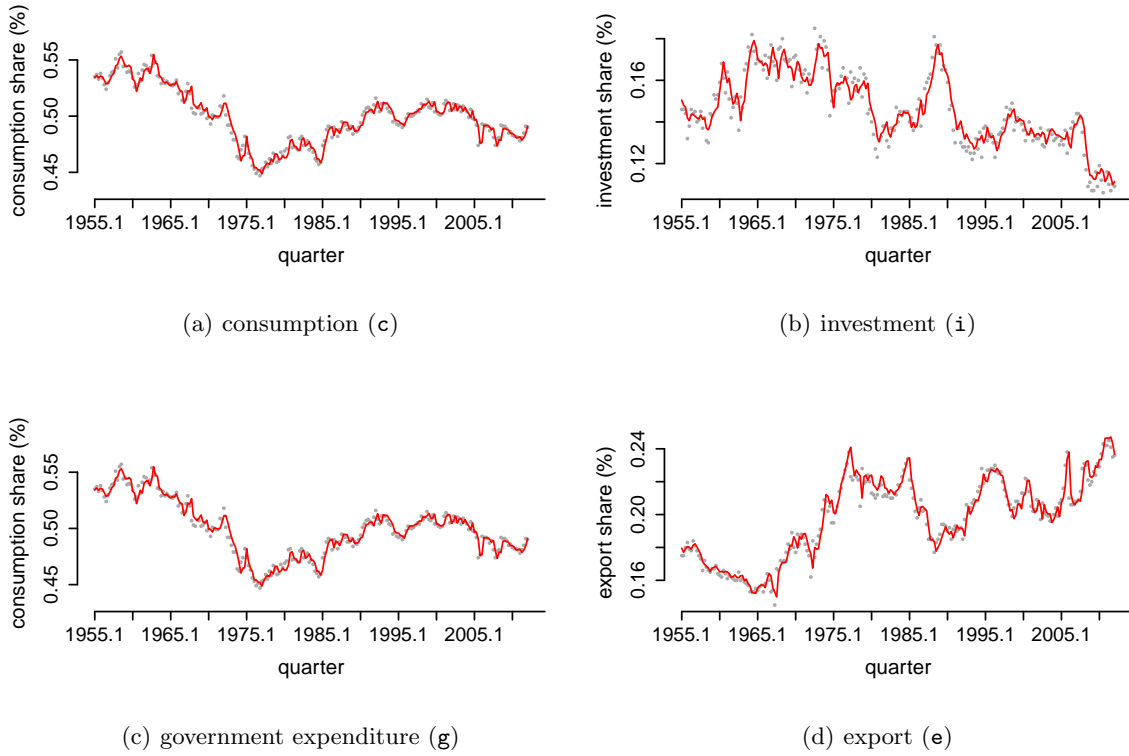


Figure 5: Filtered estimated observational mean (solid line) based on latent states posterior means for the quarterly shares of U.K. gross final expenditure for the period from 1955.1 to 2012.3. The points represent the true data set.

The smoothed posterior mean estimates ($E[\mu_t|D_T]$) for all data series are represented in Figure 7. Although we treated each time series separately the estimates obtained are consistent, in the sense that, at each instant, the sum of the estimated means are approximately one. This behaviour indicates that,

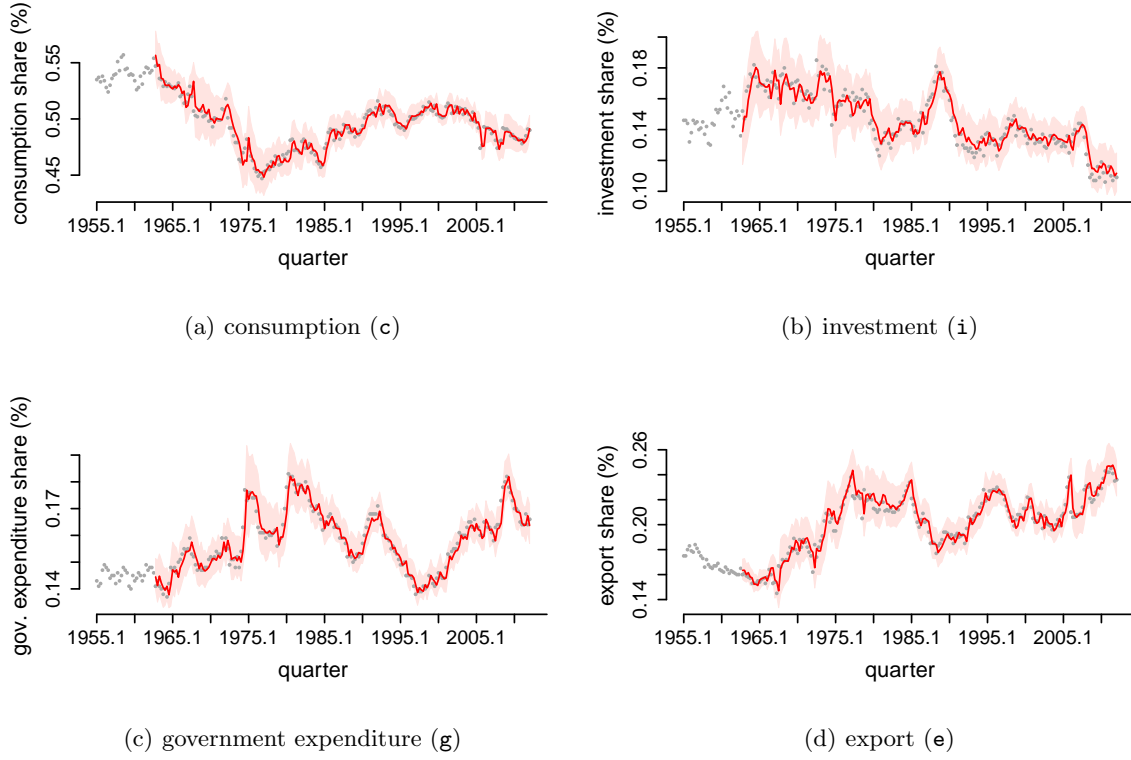


Figure 6: One-step-ahead prediction for the quarterly shares of U.K. gross final expenditure for the period from 1963.2 to 2012.3. The solid line represents the predictive distribution means ($E[y_t|D_{t-1}]$) and the shaded region the 95% HPD predictive credibility intervals. The points represent the true data set in each case.

despite the simplicity of the model used in this application, the behaviour of the series is well captured by the proposed model.

A subsets of the data set used in this application have already been analyzed by Mills (2010). Under a classical point of view, Mills (2010) estimated an order 2 VAR model, using a multivariate normal distribution to model a transformation of the original data as

$$\log\left(\frac{c}{e}\right), \quad \log\left(\frac{i}{e}\right) \quad \text{and} \quad \log\left(\frac{g}{e}\right), \quad (18)$$

where c , i , g and e represent consumption, investment, government expenditure and export rates, respectively.

In order to ascertain whether there is any advantage in analysing the data in their original scale we reanalysed these data set transforming them as proposed by Mills (2010) (according to equations (18)), replacing the observational beta distributions with univariate normal distributions for each series. Again

investment

quarter (t)	y_t	mean	mode	IC _{95%}
2011.1	0.106	0.117	0.117	[0.105 , 0.131]
2011.2	0.112	0.112	0.111	[0.098 , 0.126]
2011.3	0.116	0.113	0.113	[0.101 , 0.126]
2011.4	0.110	0.116	0.116	[0.103 , 0.130]
2012.1	0.107	0.114	0.114	[0.102 , 0.127]
2012.2	0.110	0.110	0.110	[0.097 , 0.124]
2012.2	0.109	0.112	0.111	[0.100 , 0.125]

Table 3: Point and interval estimates of the quarterly investment (i) shares of U.K. gross final expenditure based on the predictive distributions $p(y_t|D_{t-1})$.

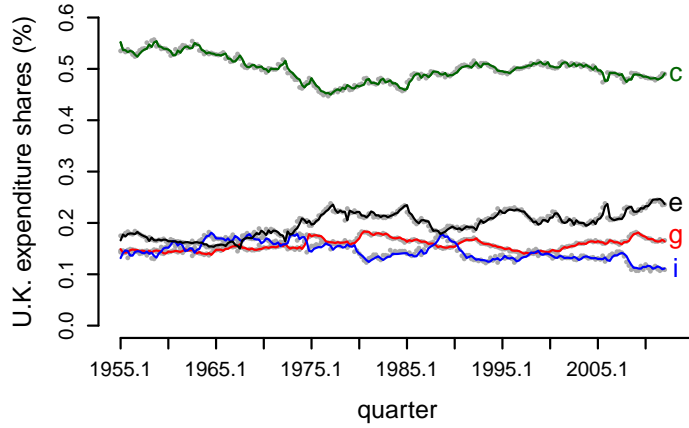


Figure 7: Smoothed posterior estimates for the observational means for quarterly consumption (c), investment (i), government expenditure (g) and export (e) shares concerning expenditure in the UK economy over the period 1955.1 to 2012.3. Gray continuous lines represent true observations.

we chose to model each series separately using analogous structures to those adopted in the beta case and assuming different discount factors for cases that include dynamics for the latent variables. According to the model comparison criteria used in this article, the best fitted standard model was the one in which we assumed a second-order VAR model for the observational means and a first-order model for the precisions,

	VAR(2)Pol(1) beta model			VAR(2)Pol(1) normal model		
	MSE	LL	LPD	MSE	LL	LPD
consumption	0.400e−4	759.776	725.504	0.061	472.727	474.535
investment	0.505e−4	730.774	698.827	0.112	591.338	374.790
gov. expenditure	0.101e−4	899.035	862.609	0.101	591.906	440.599

Table 4: Mean square error (MSE) based on one-step-ahead forecasting, joint log-likelihood (LL) and log-observed predictive density (LPD) based on consumption, investment, government expenditure and export rates for the period from 1963.2 to 2012.3, obtained from different models.

assuming a discount factor equal to 0.90 to specify the error evolution covariance matrices of the latent variables associated with precision structure.

To compare the performance of the best beta model with the corresponding normal one (both with Var(2)Pol(1)), we recalculated the normal model fit measures correcting each measure through the Jacobian of the transformation, in order to obtain adjustment measures in a same scale. The results for the fit measures for the different models can be seen in Table 4. Its possible to see that all the criteria that take into account one-step-ahead predictive distribution estimates of each of the series indicate a better performance of the beta model. Indeed, for the three considered series, the beta model had lower MSE and higher LL and LPD for all cases, giving evidence that the modelling of the data in their original scale has advantages regarding the predictive ability of the model.

5 Conclusions and Additional Comments

In this paper we propose a method for estimation and prediction of dynamic models whose observations follow distributions of the two-parameter exponential family. The estimation in the proposed partially specified model class, represented by equations (3) and (4), is based on a extension of the conjugate updating algorithm of West et al. (1985). The main idea of this new method is to explore properties of conjugacy in the exponential family and linear Bayes estimation, allowing the quick updating of both mean and precision model parameters through analytical strategies, avoiding computationally intensive methods such as those based on Monte Carlo estimation.

Our algorithm stands out mainly for two reasons: first it treats a very general class of models with observations in the exponential family, which allows modelling data in their original scale, such as in

McCullagh and Nelder (1989)' MLG. Second, the introduction of a second link function in the model allows treatment of overdispersion and heteroscedasticity in data, and allows the precision structure of the model to be dynamically treated, efficiently capturing the data behaviour even through the use of partially specified models.

Simulated studies presented by Souza (2013), assuming different observational models in the two-parameter exponential family, show that the proposed method generated satisfactory results both as regards obtaining point and interval estimates for the parameters, as in steps-ahead forecasting. The applications to real data presented in Section 4 of this paper also illustrate the good performance of the proposed algorithm and demonstrate the relevance of modelling data in their original scale.

Although use of MMG has been shown to be a good alternative to reduce the dimensionality of the system treated in Section 3.2, we intend to study other alternatives for reducing the system (5). Also with respect to the use of the generalized method of moments, we intend to study the choice of weights matrix $\mathbf{\Omega}_k$ with the aim of checking whether there is any gain in quality of estimates by introducing an iterative choice of weights matrix $\mathbf{\Omega}_k$, as discussed in Newey (1993) and Hamilton (1994).

As the main extension of this work we intend to extend the conjugate updating algorithm in order to treat classes of multi-parameter and multivariate models, such as models whose observations follow Dirichlet or multinomial distributions, the parameters of which can be explained by different link functions.

References

- Banerjee, A. K. and Bhattacharyya, G. K. (1979) Bayesian Results for the Inverse Gaussian Distribution with Application. *Technometrics*, **21**, 247–251.
- Bernardo, J. M. and Smith, A. F. M. (1994) *Bayesian Theory*. John Wiley & Sons.
- Cepeda, E. C. and Gamerman, D. (2005) Bayesian methodology for modeling parameters in the two parameter exponential family. *Revista Estadística*, **57**, 93–105.
- Cepeda, E. C., Migon, H. S., Achcar, J. A. and Garrido, L. (2011) Generalized linear models with random effects in the two parametric exponential family. *Tech. rep.*, Universidade Federal do Rio de Janeiro.
- Dey, D. K., Gelfand, A. E. and Peng, F. (1997) Overdispersed generalized linear models. *Journal of Statistical Planning and Inference*, **64**, 93–107.
- Efron, B. (1986) Double exponential families and their use in generalized linear regression. *Journal of the American Statistical Association*, **81**, 709–721.

- Gelfand, A. and Dalal, S. (1990) A note on overdispersed exponential families. *Biometrika*, **77**, 55–64.
- Hamilton, J. (1994) *Time series analysis*. Princeton University Press.
- Lindsay, B. G. (1986) Exponential family mixture models. *The Annals of Statistics*, **14**, 124–137.
- McCullagh, P. and Nelder, J. A. (1989) *Generalized linear models*. Chapman & Hall.
- Mills, T. (2010) Forecasting compositional time series. *Qual Quant*, **44**, 673–690.
- Nelder, J. A. and Lee, Y. (2001) Generalized linear models for the analysis of taguchi-type experiments. *J. Probab. Stat.*, **14**, 207–221.
- Nelder, J. A. and Wedderburn, R. W. M. (1972) Generalized linear models. *Journal of the Royal Statistical Society. Series A (General)*, **135**, 370–384.
- Newey, W. (1993) Efficient estimation of models with conditional moment restrictions. *Handbook of Statistics*, **11**, 419–454.
- Poole, D. and Raftery, A. E. (2000) Inference for Deterministic Simulation Models: The Bayesian Melding Approach. *Journal of the American Statistical Association*, **95**, 1244–1255.
- da Silva, C. Q., Migon, H. S. and Correia, L. T. (2011) Dynamic Bayesian beta models. *Computational Statistics & Data Analysis*, **55**, 2074–2089.
- Souza, M. A. O. (2013) *Aproximações analíticas e inferência em modelos na família exponencial biparamétrica*. Ph.D. thesis, Universidade Federal do Rio de Janeiro.
- Team, R. D. C. (2011) *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. URL <http://www.R-project.org/>.
- West, M. and Harrison, J. (1997) *Bayesian Forecasting and Dynamic Models*. Springer Verlag.
- West, M., Harrison, P. and Migon, H. (1985) Dynamic generalized linear models and Bayesian forecasting (with discussion). *Journal of the American Statistical Association*, **80**, 7397.
- Yin, G. (2009) Bayesian Generalized Method of Moments. *Bayesian Analysis*, **4**, 191–208.